

A Security Monitoring Framework For Virtualization Based HEP Infrastructures

A. Gomez Ramirez¹, M. Martinez Pedreira², C. Grigoras², L. Betev², C. Lara¹ and U. Kebschull¹ for the ALICE Collaboration

¹Infrastructure and Computer Systems for Data Processing (IRI), Goethe-University Frankfurt

²CERN, Geneva, Switzerland

E-mail: andres.gomez@cern.ch

Abstract. High Energy Physics (HEP) distributed computing infrastructures require automatic tools to monitor, analyze and react to potential security incidents. These tools should collect and inspect data such as resource consumption, logs and sequence of system calls for detecting anomalies that indicate the presence of a malicious agent. They should also be able to perform automated reactions to attacks without administrator intervention. We describe a novel framework that accomplishes these requirements, with a proof of concept implementation for the ALICE experiment at CERN. We show how we achieve a fully virtualized environment that improves the security by isolating services and Jobs without a significant performance impact. We also describe a collected dataset for Machine Learning based Intrusion Prevention and Detection Systems on Grid computing. This dataset is composed of resource consumption measurements (such as CPU, RAM and network traffic), logfiles from operating system services, and system call data collected from production Jobs running in an ALICE Grid test site and a big set of malware. This malware was collected from security research sites. Based on this dataset, we will proceed to develop Machine Learning algorithms able to detect malicious Jobs.

1. Introduction

Frequently in HEP computing, and also in general purpose Grid computing, user supplied code and data are deployed and executed in farms around the world, while the exact location is normally irrelevant. This allows scientists from many areas beyond physics to use huge computational power to solve complicated scientific problems, such as weather modeling, brain simulation, among others. However it also creates cyber-security challenges. Operators and administrators need tools to monitor for security incidents. User code and data should be isolated from different users, also from the physical computers and networks, in order to restrict access to sensitive elements in the organizations.

We propose a novel paradigm and developed a framework that focus on protecting and monitoring user payload execution. Moreover, it enforces isolation in the environment in such a way that Jobs cannot access sensitive resources. This tool enables the Job behavior analysis in order to detect possible intrusions. This is accomplished by collecting and processing data generated by Jobs such as logs, system calls and resource consumption data. Traditional Intrusion Detection and Prevention Systems (IDPS) perform attack detection by using fixed rules based on signatures, identical to traditional monitoring systems. Therefore, we employ Machine Learning (ML) to overcome the mentioned drawbacks, achieving generalization among

attack variants. Currently there is no tool that provides isolation, while monitoring security incidents by ML algorithms in Grid computing [1].

The authors have defined a threat model that guides the design and implementation of the described framework [1]. It is devised to detect attackers in the protected system trying actions like the following:

- Exploit unknown or unfixed software/hardware vulnerabilities.
- Listen to user network traffic to gather sensitive clear text information.
- Perform a 'man in the middle' attack.
- Tamper other user Jobs.
- Escalate privileges.
- Access sensitive server configuration data.

As a proof of concept we are implementing the described framework for the ALICE Grid at CERN. ALICE (A Large Ion Collider Experiment) is a dedicated Pb–Pb detector designed to exploit the physics potential of nucleus-nucleus interactions at the Large Hadron Collider at CERN [2, 3]. The ALICE experiment has developed the ALICE production environment (AliEn) [4], which implements many components of the Grid technologies that are needed to analyze HEP data. Through AliEn, the computing centers that participate in the ALICE Grid can be seen and used as a single entity. Any available node executes Jobs and file access is transparent to the user, wherever in the world a file might be [5]. Figure 1 shows a picture of the ALICE Grid.

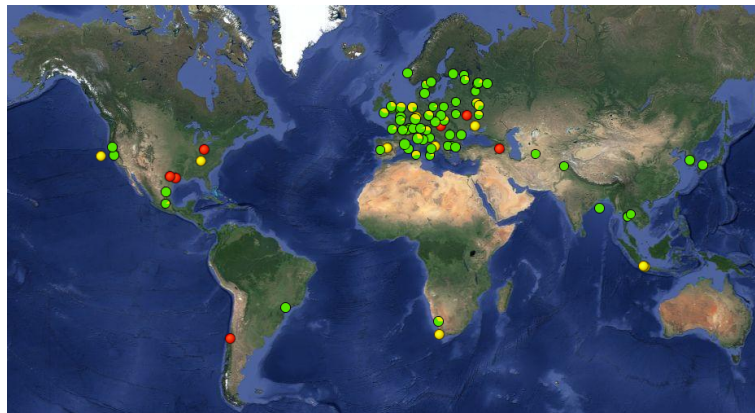


Figure 1. ALICE Grid computing farms around the world.

This document is organized as follows: section 2 introduces a security by isolation strategy for a distributed system. Section 3 explains a method for collecting relevant data from the isolated infrastructure. Section 4 shows how the collected data can be used to determine the security status of the system. In section 5 we detail our Intrusion Detection and Prevention model. Section 6 summarizes the current state of the project and the challenges faced in the design and implementation of the desired methodology. Finally, section 7 gives conclusions on the work done.

2. Security by isolation

Security by isolation enforces application space separation [8]. The idea is that, if one process is compromised and utilized to attack the entire system, other components can stay untouched.

Several technologies provide secure isolation. Virtual Machines (VM) and Linux Containers (LC) are very popular examples.

2.1. Linux containers

LC are an extension of the virtual memory concept to allow the isolation of network interfaces, the PID tree and mount points [9]. Separation of containers from the rest of the system is enforced by the Kernel, so they can not affect the host or other containers. LC technology uses namespaces and Cgroups [10] to have a private view of the system and a limited resource assignment.

Containers provide a set of advantages over VM. They are lightweight and fast, boot in milliseconds and have just a few MB of intrinsic disk and memory usage. It has been shown [11], that they provide a better performance than VMs. Commonly, VM are used in Grid and Cloud computing to achieve isolation, however LC performance and comparable security features make it a suitable alternative[6, 7].

2.2. Proposed isolation architecture

We propose the usage of LC to enforce HEP Grid site user isolation, also extensible to broader scientific computing and clouds. To achieve this we require a batch Job orchestrator allowing the execution of user processes in containers on computing clusters. Section 6 gives further details about this requirement and the selected solution. As shown in Figure 2, we switch from an environment without isolation, where Jobs have access to the server and other user Jobs, to an environment where Jobs run in their one process space, without access to other Jobs or sensitive components.




Figure 2. Desired isolation scenario.

An isolated environment for Job execution is not enough. Jobs could still perform several kind of attacks or not allowed activities such as, Distributed Denegation of Service (DDoS), Bitcoin mining, malware/botnet hosting, among others. Consequently, we need to monitor the activity and detect incidents inside the LC, as described in the next sections.

3. Monitoring data mining

HEP distributed computing systems use continuous automated monitoring to help administrators to find and fix situations affecting the normal operation [12]. The resulting monitoring data can be used to find or even predict software and hardware failures. It is a valuable source of security information as well. In this document we focus on the measurement of metrics related to the batch Jobs being submitted to the distributed system. There are several relevant metrics that we can collect, for instance:

- Job and system logs.

- System call sequence.
- Resource usage data (such as CPU, RAM and network traffic).

Furthermore our goal is to chose the best information about Job behavior without affecting the habitual performance. We decided to employ data mining and intelligent algorithms, given their ability to find correlations and analyze trends in big datasets [13], in order to provide a better understanding of security related events.

4. Machine learning based security monitoring

Machine Learning is a set of mathematical models that simulate the human learning abilities[14, 15]. In the context of Intrusion Detection, ML helps analyzing big amounts of data by learning the expected behavior and identifying abnormal situations. Traditional industrial IDS use rather fixed rules and search for known attack signatures. However they have problems when unknown or slightly different intrusion methods are used [13]. We have selected supervised training to analyze the collected data. In supervised training, a set of already classified and tagged data (training dataset) is used to model a function (for example a Neural Network) in order to make it able to classify new unseen data (test dataset) [16].




Figure 3. Monitored data gathering and processing.

4.1. Training dataset

We have collected a training and testing dataset by gathering monitoring data from production ALICE Grid Jobs (Figure 3). Additionally we have executed a big set of Linux malware samples. The data on the first part is tagged as normal data and the second as malicious. This dataset is utilized to compare several Machine Learning algorithms to find the one that gives the best accuracy. Following is the list of ML algorithms selected. They will be tested to define which one gives the best accurate results for our dataset:

- Support Vector Machines.
- Multilayer Neural Networks.
- Recurrent Neural Networks.

Figure 4 shows a scheme of the proposed architecture for the ML usage.

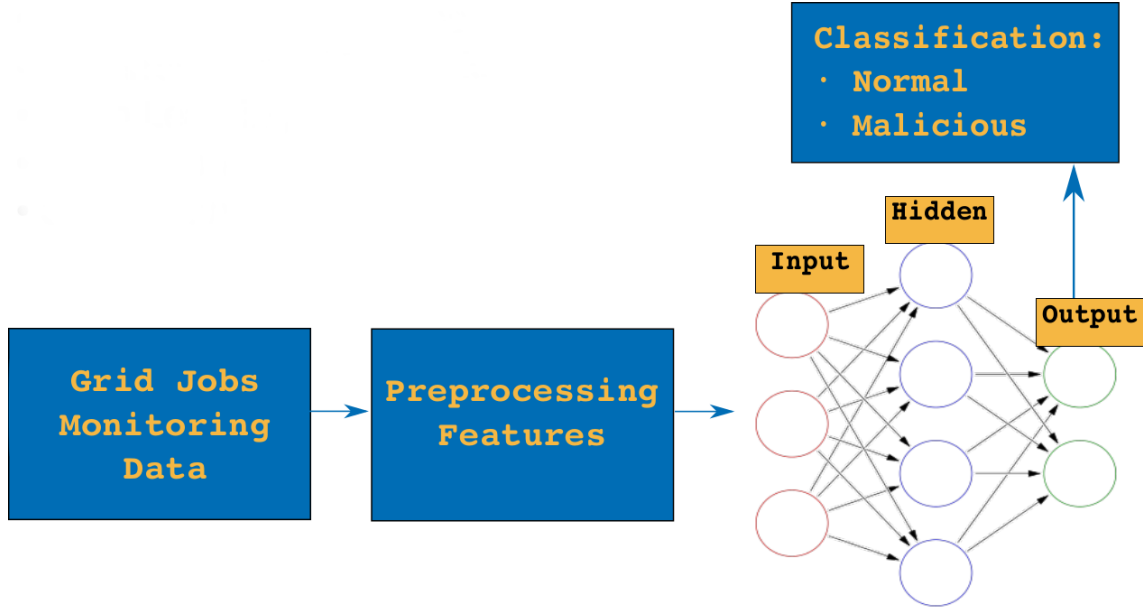


Figure 4. Proposed architecture.

5. Intrusion prevention and detection

When analyzing Job monitoring data, our goal is finding security incidents. This security incidents can be found by analyzing anomalies in the system, things that go beyond the common state, probably caused by malicious software. Besides, even if our execution environment is sandboxed, there are many possible attacks that can still affect the distributed infrastructure. If a user's Job is misbehaving, the proposed framework should raise an alarm and perform predefined actions, for instance terminate the malicious processes. Figure 5 shows the desired implementation of the proposed system regarding Intrusion Detection.



Figure 5. Proposed Intrusion Detection in the worker nodes.

5.1. Challenges

Improving the provided security should not impact the system performance. This is especially important in HEP computing. On the other hand, we also need innovative ways to analyze the trace and log data in an efficient way. Another important challenge is to reduce the amount of false positives and false negatives, since the system administrators rely on the accuracy of the security monitoring framework.

6. Proof of concept and testing environment

So far we have already deployed a testing ALICE Grid site based on AliEn [4] in a local Linux cluster, with five Ubuntu 14.04 nodes. In order to orchestrate and run the Jobs inside Linux Containers we have tested three different tools that offer such functionality:

- Kubernetes [17].
- Apache Mesos [18].
- Docker Swarm [19].

At the end we have decided to work with Docker Swarm, because it allows to carry out the simplest deployment, which is an important requirement for our research environment. We use Docker [20] as LC engine, with Centos 6 [21] container images. We have developed AliEn interfaces for the mentioned batch systems. CVMFS [22] is installed on the hosts and shared as a volume inside the AliEn container to allow access to HEP libraries. Currently we execute one Job per container. This is useful to increase the traceability between different Jobs. Also, this is the natural micro service model for LC.

As a monitoring infrastructure for collecting data from normal Grid Jobs we have Prometheus [23] and Sysdig [24]. Prometheus allows to take resource usage data directly from containers and collect it via a RESTful interface. Sysdig enables to capture system calls in Linux OS in a fast and reliable way. We have developed a custom Python library to integrate these tools and make them fit our needs. This infrastructure has been utilized for the execution and measurement of ALICE production Jobs, that are tagged as normal Jobs.

A network isolated machine was used for malware data collection. This machine has the same setup as the Grid worker nodes. We have downloaded a set of 10000 Linux malware samples from a security research web site [25]. We ran the samples and collected the same information as for the normal Jobs (logs, sequence of system calls, resource usage data). Finally we obtained a combined dataset that allows to train and test our selected Machine Learning algorithms. A representation of the implemented components are shown in the Figure 6.

7. Conclusions

Distributed computing is a fundamental component of High Energy Physics collaborations. Improving security in this kind of infrastructures requires innovative tools to automatically detect security related incidents. Security by isolation is also necessary to protect sensitive components allowing traceability on Job activities. We propose the usage of Linux Containers in order to provide isolation without highly decreasing the expected performance. We use Machine Learning techniques to provide generalization, overcoming common IDS difficulties on finding even slightly different threats. We describe the ongoing development process of a new security monitoring framework for Linux Containers based HEP infrastructures. This is being tested as a proof of concept for the ALICE experiment at CERN. We have collected a dataset of normal and malicious monitoring information from Grid Jobs and malware samples, that will be utilized to train and test ML algorithms. These algorithms should enable autonomous Intrusion Detection and Prevention as an important component of the proposed new framework. As future work we plan to explore how our approach can be used to detect anomalies that go beyond the security scope, for instance to find hardware failures or even human mistakes.

Acknowledgments

Authors acknowledge assistance from CERN security department specially Stefan Lueders and Romain Wartel. This work is supported by the German Federal Ministry of Education and Research.

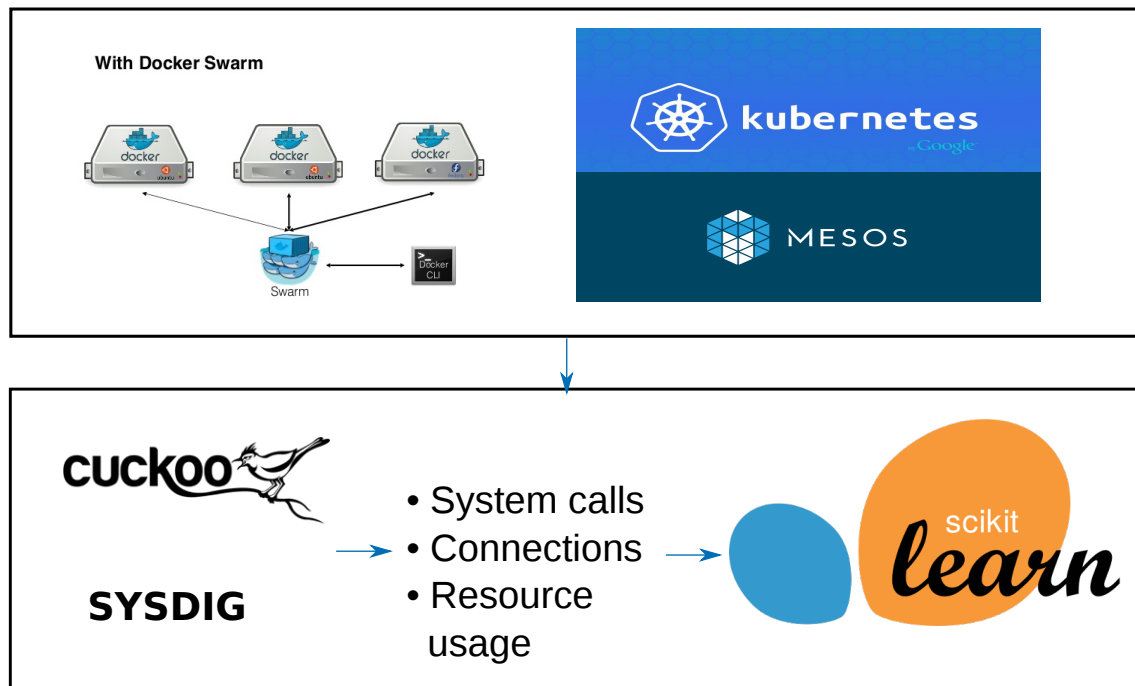


Figure 6. Proof of concept implementation.

References

- [1] Gomez Ramirez A, Lara C and Keschull U 2015 *JPCS*. Intrusion Prevention and Detection in Grid Computing - The ALICE Case.
- [2] The ALICE collaboration 1995 ALICE - Technical Proposal for A Large Ion Collider Experiment at the CERN LHC, CERN, Geneva, Rep. CERN-LHCC-95-71; LHCC-P-3.
- [3] The ALICE collaboration 2008 *JINST*. The ALICE Experiment at the CERN LHC, , vol. 3, no. 08, Aug.
- [4] Begnasco S et al. 2015 *JPCS*. AliEn: ALICE environment on the GRID. IOP Publishing.
- [5] Christoph E et al. 2005 *Technical Design Report LCG*. LHC computing Grid: Technical Design Report. Version 1.06. 20 Jun 2005. CERN, Geneva. Available from: <https://cds.cern.ch/record/840543>.
- [6] Abed A, Clancy C and Levy C 2015 *Security and Trust Management, Lecture Notes in Computer Science*. Intrusion Detection System for Applications Using Linux Containers. Springer International Publishing.
- [7] Modi C et al. 2013 *JNCA*. A survey of intrusion detection techniques in Cloud. Jan;36(1):42-57.
- [8] Mansfield-Devine S 2010 *Computer Fraud & Security*. Security through isolation. 5, 8 - 11. 1361-3723.
- [9] Graber S, 2014 Ubuntu Foundations Team, LXC 1.0: Blog post series. [updated 2014 Jan 17; cited 2015 May 11]. Available from: <https://www.stgraber.org/2013/12/20/lxc-1-0-blog-post-series/>
- [10] Menage P 2007 *Proceedings of the Linux Symposium*. Adding Generic Process Containers to the Linux Kernel. June 27.
- [11] Shiseki A 2016 *The 22nd International Conference on Computing in High Energy and Nuclear Physics, CHEP 2016*. Cloud Computing as a Scientific and Technical Application Development and Execution Platform. October 10-14.
- [12] Lara C et al. 2011 *JPCS*. Autonomous System Management for the ALICE High-Level-Trigger Cluster using the SysMES Framework. doi:10.1088/1742-6596/331/5/052003.
- [13] Azad V 2013 *IJITCS*. Data Mining in Intrusion Detection: A Comparative Study of Methods, Types and Data Sets. 75-90.
- [14] Yiping L et al. 2012 *IJACT*. An Intrusion Detection Approach Using SVM and Multiple Kernel Method. Jan 31;4(1):463-9.
- [15] Gascon H et al. 2014 *ACM Press; 2013*. Structural detection of android malware using embedded call graphs. [cited 2014 Nov 17]. p. 45-54. Available from: <http://dl.acm.org/citation.cfm?doid=2517312.2517315>
- [16] Bishop C 2006 *Library of Congress ISBN* Pattern recognition and machine learning. New York. Springer.
- [17] Google, Kubernetes. 2016 [updated 2016 Dec 1; cited 2016 Dec.1]. Available from: <http://kubernetes.io/>

- [18] Apache foundation, Mesos 2016 [updated 2016 Dec 1; cited 2016 Dec.1]. Available from: <https://mesos.apache.org/>
- [19] Docker, Docker Swarm 2016 [updated 2016 Dec 1; cited 2016 Dec.1]. Available from: <https://docs.docker.com/swarm/>
- [20] Docker, Docker 2016 [updated 2016 Dec 1; cited 2016 Dec.1]. Available from: <https://docs.docker.com>
- [21] CentOS Project, Centos6 2016 [updated 2016 Dec 1; cited 2016 Dec.1]. Available from: <https://www.centos.org/>
- [22] Buncic P et al 2008 *Proceedings of the XII. International Workshop on Advanced Computing and Analysis Techniques in Physics Research (ACAT08)*. CernVM - a virtual appliance for LHC applications. Erice, PoS(ACAT08)012.
- [23] Prometheus 2016 [updated 2016 Dec 1; cited 2016 Dec.1]. Available from: <https://prometheus.io/>
- [24] Draios, Sysdig 2016 [updated 2016 Dec 1; cited 2016 Dec.1]. Available from: <http://www.sysdig.org/>
- [25] VirusShare 2016 [updated 2016 Dec 1; cited 2016 Dec.1]. Available from: <https://virusshare.com/about.4n6>